REPORT DOCUMENTATION PAGE

AD-A201 458	**1b. RESTRICTIVE MARKINGS**
ULE OCT 2 0 1988	**3. DISTRIBUTION / AVAILABILITY OF REPORT** Approved for public release; distribution unlimited.

4. PERFORMING ORGANIZATION REPORT NUMBER(S)	5. MONITORING ORGANIZATION REPORT NUMBER(S)
	ARO 22648.16-EL

6a. NAME OF PERFORMING ORGANIZATION	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION
University of Maryland		U. S. Army Research Office

6c. ADDRESS (City, State, and ZIP Code)	7b. ADDRESS (City, State, and ZIP Code)
College Park, MD 20742	P. O. Box 12211 Research Triangle Park, NC 27709-2211

8a. NAME OF FUNDING / SPONSORING ORGANIZATION	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER
U. S. Army Research Office		DAAG29-85-K-0177

8c. ADDRESS (City, State, and ZIP Code)	10. SOURCE OF FUNDING NUMBERS			
	PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT ACCESSION NO
P. O. Box 12211 Research Triangle Park, NC 27709-2211				

11. TITLE (Include Security Classification)

Non-Monotonic Reasoning, Belief Systems, and Parallelism

12. PERSONAL AUTHOR(S)
Jack Minker and Donald Perlis

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Year, Month, Day)	15. PAGE COUNT
Final	FROM 7/1/85 TO 6/30/88	Jul 88	10

16. SUPPLEMENTARY NOTATION The view, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation

17.	COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	Non-Monotonic Reasoning, Belief Systems, Parallelism,
			Uncertain Data, Deductive Databases

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

The research can be categorized in terms of the following inter-related concepts:

• Non-monotonic reasoning and circumscription

• Belief systems and deductive databases (in which we take account of models of an intelligent agent's belief set, especially with regard to intelligent answers to queries in deductive databases)

• Resource-limited non-monotonic belief systems (a special case of the above, in which the non-monotonic nature of beliefs is studied)

• Parallelism in non-monotonic belief systems (in which parallel methods of computation are exploited with regard to the above)

20. DISTRIBUTION / AVAILABILITY OF ABSTRACT	21. ABSTRACT SECURITY CLASSIFICATION
☐ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT. ☐ DTIC USERS	Unclassified

22a. NAME OF RESPONSIBLE INDIVIDUAL	22b. TELEPHONE (Include Area Code)	22c. OFFICE SYMBOL

DD FORM 1473, 84 MAR   83 APR edition may be used until exhausted.   SECURITY CLASSIFICATION OF THIS PAGE

All other editions are obsolete.   UNCLASSIFIED

1

Non-Monotonic Reasoning, Belief Systems, and Parallelism

FINAL REPORT

Jack Minker and Donald Perlis

U. S. ARMY RESEARCH OFFICE

ARO PROPOSAL NUMBER: 22648-EL

CONTRACT OR GRANT NUMBER: DAAG29-85-K-0177

PERIOD COVERED BY REPORT: 1 July 1985 - 30 June 1988
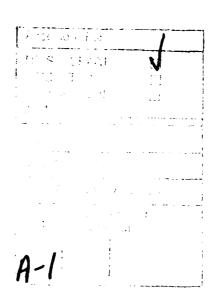
NAME OF INSTITUTION: University of Maryland

APPROVED FOR PUBLIC RELEASE:

DISTRIBUTION UNLIMITED.

TABLE OF CONTENTS

A-1

A. STATEMENT OF THE PROBLEM STUDIED

In this section we summarize our work over the three year period of the research grant. Our work has centered on reasoning with incomplete or uncertain data, especially with attention to areas in which parallelism could be exploited. The research contributions from this work range from theoretical investigations to experimental and practical results. Numbered references in brackets (e.g., [5]) are to the 18 papers we published during the grant period, listed in section C below; other references are found after that, and indicated by letters in brackets (e.g., [B]).

Non-monotonic reasoning (NMR) is a form of reasoning with highly unusual formal properties. A key idea in NMR is that of drawing conclusions about the typicality of an entity x with respect to a property P, so that, when no contrary evidence is present, Px is concluded. That is, in such circumstances x is deemed to be typical (to have property P). For instance, if Soviet missiles are typically in good working order, and m37 is a given Soviet missile, then in the absence of evidence to the contrary it will be supposed that m37 is in good working order. Since rarely are all the relevant facts in a complex setting known, such reliance on typicalities in terms of available evidence becomes a matter of considerable importance.

Some of our efforts have gone into studying theoretical aspects of NMR, in the specific area of circumscription. Specifically, the phrase "when no contrary evidence is present" given above is very difficult to incorporate into a deductive formalism. We have considered this from several directions, including circumscription and step logic. However, NMR also (and principally) has applications to the study of intelligent (commonsense) reasoning. such as the study of belief systems of intelligent agents. For instance, communicating agents tend to make significant use of tacit cooperation principles, such as giving not merely correct answers to questions, but also *informative* and *not misleading* answers. This in turn depends strongly on the set of beliefs (and intentions) of the agents involved. NMR can be used to model such principles via making suppositions about what another agent does or does not know or believe, again based on ideas of typicality. Thus ths study of belief systems forms an important area of research involving NMR.

Moreover, questions of efficiency plague NMR, since many of the formal methods used involve either semi-decidable or undecidable problems, and also since realistic databases for commonsense are so large that even ordinary (monotonic) reasoning is slow. For these reasons. among others, many researchers in AI look to both resource-limited and parallel inference as a way to significantly speed up deductions. We have been exploring both parallel and resource-limited methods for calculating non-monotonic conclusions.

Thus our research can be categorized in terms of the following inter-related concepts:

• Non-monotonic reasoning and circumscription

• Belief systems and deductive databases (in which we take account of models of an intelligent agent's belief set, especially with regard to intelligent answers to queries in deductive databases)

• Resource-limited non-monotonic belief systems (a special case of the above, in which the non-monotonic nature of beliefs is studied)

• Parallelism in non-monotonic belief systems (in which parallel methods of computation are exploited with regard to the above)

During the three year grant period, we have published 11 conference papers and 7 journal articles. One journal article [17] was an invited paper (and also an invited conference talk). and one paper [6] received honorable mention in a conference.

B. SUMMARY OF THE MOST IMPORTANT RESULTS

1. Non-monotonic reasoning and circumscription.

McCarthy [A] has developed a form of non-monotonic reasoning called circumscription. Circumscribing a predicate P is an important aspect of most problems in artificial intelligence in which the user generally wants to consider only the facts at hand and no others. Given a sentence, A(P), of the first order predicate calculus, in which the predicate P appears, then McCarthy's original schema may be written as,

$$C(Z): A(Z)\&(x)(Z(x) \to P(x)) \to (y)(P(y) \to Z(y)).$$

That is, if for any formula $Z(x)$, it can be shown that the first two parts of the schema are true, then one can conclude that only Z-things may be P-things. Thus if missiles m7 and m8 are known to be functional, m9 to be non-functional, and no data is available on m10 and m11, then ordinary circumscription of the predicate Unfunctional(x) would lead to the conclusion that m10 and m11 are not Unfunctional, i.e., that they are functional. Here P(x) would be Unfunctional(x), and $Z(x)$ could be taken to be $x = m7 \lor x = m8 \lor x = m9$.

We generalized McCarthy's schema to allow protecting certain atoms P(a) from being decided negatively. Suppose as before that missiles m7 and m8 are known to be functional, m9 to be non-functional, and no data is available on m10 and m11. However, suppose also that we have some evidence that m10 may be Unfunctional, but it is inconclusive as yet. We may wish to block the default conclusion that m10 is functional, and await further evidence. In this case McCarthy's schema, as it has been written, does not apply. The schema was revised by us as

$$C(Z): A(Z)\&(x)(Z(x)/S(x) \to P(x)) \to (y)(P(y)/S(y) \to Z(y)),$$

where by $Z(x)/S(x)$ we mean $Z(x)\&\neg S(x)$.

This concept of "protected circumscription" is described in our paper [2]. McCarthy [B] has also extended his earlier concept of circumscription, and his extension subsumes the concept of protected circumscription, although as noted below it does not lend itself as readily to computation in certain cases.

In our paper we study computational issues and show how protected circumscription can be incorporated into PROLOG by a modification to the negation operator. We also investigated computing circumscription in the case of simple forms of data (called Horn data) with additional protection (indefinite data), an intermediate investigation between Reiter's result [C] on predicate completion and Lifschitz's efforts [D] to make general (formula) circumscription more efficient as a computational tool. Reiter's work showed a close tie between McCarthy's circumscription and Clark's predicate completion [E]. We investigated a similar tie between an extended version of circumscription involving protected data, and an extended version of predicate completion. When we have a fully ground atomic protected theory, we show that an extension to the relational algebra can be used to obtain all (and only) correct answers. When general Horn axioms are added to the protected theory, we show that Horn axioms also can be used to compute sound answers; however, some correct answers will not be found.

In another study, we investigated the model theory of the notion of circumscription, and found completeness theorems that provide a partial converse to a result of McCarthy. We show that the circumscriptive theorems are precisely the truths of the minimal models (i.e., models in which the P-things form subsets of the domain which cannot be made smaller and remain consistent with the axioms) in the case of various classes of theories, and for various versions of circumscription. This of course is what is expected, since minimization of P-things is the purpose of circumscribing the predicate P. However, Etherington et al. [F] and Kueker

[G] have noted that this is not always the case, so that it is reassuring to have such a result providing intuitive correctness. This work we reported in [3].

Sets can play an important role in circumscription's ability to deal in a general way with certain aspects of commonsense reasoning. Kueker's result mentioned above indicates that sentences that intuitively one would want circumscription to prove, are nonetheless not so provable in a formal setting devoid of sets. Furthermore, when sets are introduced, first-order circumscription handles these cases very easily, obviating the need for second-order circumscription. The *Aussonderungs* axiom of ZF set theory plays an intuitive role in this shift back to a first-order language. The paper [10] presents this work. Further ramifications of set theory in commonsense reasoning are presented in [9], where it is argued that set theory provides a powerful addition to commonsense reasoning, facilitating expression of meta-knowledge, names, and self-reference. Difficulties in establishing a suitable language to include sets for such purposes are discussed, as well as what appear to be promising solutions. One issue highlighted in this paper is that it is important to be able to regard the universe of discourse itself as being an object of discourse; this runs into problems related to Russell's paradox. Possible solutions are suggested.

2. Belief systems and deductive databases

Our work on belief systems includes much that also deals with non-monotonic reasoning. In this section we describe that portion which emphasizes belief systems as essential aspects of any intelligent attempt to answer queries in deductive databases. This includes what is variously described as "cooperative query answering" and the "cooperation principle." The idea is that often a literal answer to a query (or question) is not what is wanted, or is even misleading. For instance, the query "Did Smith take EFL 620?" may have the literal answer "No." But the more informative answer may be "There is no EFL 620" if there is no such course as EFL 620.

Gal and Minker have in their papers [1,13] pursued natural language aspects to the query-answer issue. They have developed a general theory as to how to provide natural language intelligent or cooperative (rather than simply literal) responses to queries. Their approach takes advantage of integrity constraints that exist for a database. Detailed heuristics have been developed for this purpose. We are currently considering an implementation of our general method in PROLOG. As a first step in this regard we have developed a meta-interpreter in PROLOG [18] that combines integrity constraints with axioms in logic programs. The meta-interpreter is currently running and we plan to extend the meta-interpreter to permit cooperative answers to be obtained.

Miller and Perlis [12,14] studied interactions between theorem-proving and natural language processing. In particular, they discuss Ohlbach's claim that first-order logic is not well suited to handling certain problems involving truth of utterances. They provide another interpretation of the problem using indexicals, and axiomatize it so that the desired results follow. They conclude by suggesting a broader context for dealing with utterances in automatic theorem-proving.

Another issue involving language is that of meaning, in the following sense: If a system employs symbols, in what sense are they symbolic, of *what* are they symbolic, and in what sense is it the *system* that makes them symbolic? In other words, what does it take for a system to be such that it can have beliefs that are *about* something? We suggest an answer based on the idea of quotation or reification, namely that both the symbol and its purported meaning must be represented within the reasoning system, for instance both the word 'dog' and some (other) representation taken for a dog (such as a picture). Here the word can be regarded as a linguistic reification of the picture. Such a scenario might be able to account for the drawing of conclusions as to the meaning of a word. This is reported in [11].

In [17] Minker discusses work in deductive databases that has taken place over the last 30 years. He presented an invited talk at the Principles of Databases Conference and submitted an invited paper on the subject that appeared in the Journal of Logic Programming [17]. Deductive databases are important and provide the basis for knowledge-based, belief-based, and expert database systems, since it is usually the inferences or deductions that may be drawn from given beliefs that is of interest in reasoning about intelligent behavior.

9. Resource-limited non-monotonic belief systems

Non-monotonicity is pervasive in commonsense reasoning, and therefore comes into most of our work, including that on circumscription and deductive databases above However, certain of our efforts have focussed on classical problems in the literature of non-monotonic belief systems. such as "When is it reasonable for an intelligent agent to believe a bird can fly" and "When does lack of a belief that P is true, imply the falsity of P?" We have studied a real-time mechanism, which we call step-logic, to address these problems.

We proposed that a new kind of logical study is appropriate to agents engaged in commonsense reasoning, namely, one that focuses on the steps of reasoning at any given time rather than the collection of all conclusions ever reached. This was reported in [4] by Drapkin and Perlis. We then carried out this program for the propositional case, and in particular gave a result on completeness for reasoning about agents; see [7]

The kind of resource limitation that is most evident in commonsense reasoning (i.e., in reasoning about and within a real environment) is simply the passage of time while the reasoner reasons. There is not necessarily (or even likely) any fixed and final set of consequences with which such a reasoning agent ends up. What is of interest for such an agent is not its "ultimate" set of conclusions, but rather its changing set of conclusions over time. Indeed, there will be, in general, no ultimate or limiting set of conclusions.

In contrast to already existent approaches, we have proposed step-logic to model reasoning that focuses on the on-going process of deduction. The reasoner starts out with an initial set of axioms at time 0. At some time, i, it concludes theorem α, at some later time, j, it comes up with β. Of course, this much might be said of any deductive logic. However in step-logic these time parameters can figure in the on-going reasoning itself.

An interesting problem, related by Moore, is the following. One is able to reason, "Since I don't know I have a brother, I must not." This problem can be broken down into two: the first requires that the reasoner be able to decide he doesn't know he has a brother; the second that, on that basis, he, in fact, does not have a brother (from modus ponens and the assumption that "If I had a brother, I'd know it.") The first of these seems to lend itself readily to step-logic, in that the negative reflection problem (determining when something is not known) should reduce to a simple look-up. In our paper we presented a real-time solution to Moore's Brother Problem, with computer-generated results for three different scenarios involving determining whether or not the proposition that a brother exists, holds.

Specifically, we have implemented a real-time inference mechanism that correctly infers a lack of knowledge of certain wffs; that correctly will not infer a lack when the knowledge is in fact present; and that correctly resolves a contradiction when timing is such that new knowledge arises conflicting with a prior (or simultaneous) conclusion of its lack. Of course, we have shown this only in a very limited context so far. One interesting feature is that it is not at all critical whether a contradiction is instantly resolved. In step-logic, the agent can continue reasoning step-by-step in the presence of a contradiction. The possible "spread" of invalid conclusions from a contradiction itself goes only step-by-step, thus presenting the possibility of controlling it by effective means.

Further work in non-monotonic belief systems includes a philosophical treatment in which non-monotonic (or "default") reasoning is analyzed as consisting (implicitly) of at least three further aspects, that we call oracles, jumps, and fixes, which in turn are related to the

notion of a belief. Beliefs are analyzed in terms of their use in a reasoning agent, and an idea of David Israel is embellished to show that certain desiderata regarding these aspects of default reasoning lead to inconsistent belief sets. As a consequence the handling of inconsistencies must be taken as central to commonsense reasoning, lending support for our above work in step logic. Finally, these results are applied to standard cases of default reasoning formalisms in the literature (circumscription, default logic, and non-monotonic logic), where it turns out that even weaker hypotheses lead to failure to achieve commonsense default conclusions. This was reported in [8].

One surprising result was that a "counter-example" axiom, to the effect that defaults (typicality rules) have exceptions, undoes the effectiveness of certain formal approaches to NMR. For instance, the axiom (x)(BIRD x &¬FLIES x) prevents circumscription (at least in the usual straightforward uses) from concluding FLIES tweety given the information BIRD tweety. Thus when formal reasoners are made to know certain aspects of their own forms of reasoning (e.g., that certain information is tentative and error-prone) unexpected behaviors can result. Further results on inherent difficulties in formal approaches that do not incorporate appropriate distinctions of temporal aspects of reasoning are given in [16].

4. Parallelism in non-monotonic belief systems

Parallelism has recently gained widespread attention within AI, and more generally within Computer Science. The principal motivation is the expected speed-up when many processors execute simultaneously, as opposed to one processor executing activities in sequence. Another motivation stems from the inherently parallel features of computer networks, and a third is the highly parallel processing in the human brain which is a major source of insight in cognitively-oriented studies within AI.

As another step in our efforts toward the study of real-time monitoring of the inferential process in reasoning systems, we have studied a meta-level method of representing knowledge for parallel default reasoning. A meta-level implementation that permits effective monitoring of the deductive process as it proceeds, providing information on the state of the answer procurement process, has been designed for the Parallel Inference System (PRISM) at the University of Maryland.

PRISM is a parallel logic programming system developed to execute logic programs by solving the implicit AND/OR tree that is defined by a query and a program. Parallelism is achieved by distributing portions of the tree to different processors. Although PRISM can exploit parallelism transparently the user may annotate his program and explicitly control the parallel execution. A PRISM program consists of a finite set of annotated Horn clauses defining predicates. A predicate could be defined in terms of other predicates or facts (unit clauses). Also, there is a set of built-in (predefined) predicates in the system such as SUM(X,Y,Z), MUL(X,Y,Z), X- Y, etc., that can be used in these definitions. There are two implementations of the system. One is running on McMob (a parallel computer designed and built at the University of maryland), the second is running on a Butterfly but this is a primitive version. These two implementations use message passing techniques for communication between processors.

As part of the previous ARO grant research, we devised an implementation in PRO-LOG, to be incorporated into PRISM, of a learning feature used to calculate, for purposes of issuing default answers, the current depth of inference for a query from that obtained from similar queries posed earlier. See our paper [5]. This work was our initial step toward using PRISM for our work on belief systems. Additional capabilities are required in PRISM to implement the work in [5].

C. RESEARCH PUBLICATIONS COMPLETED DURING THE GRANT PERIOD.

(1)   Gal, A. and Minker, J. "A natural language database interface that provides cooperative answers", *Proceedings, 2nd Conference on Artificial Intelligence Applications*, Dec. 11-13, 1985, Florida.

(2)   Minker, J. and Perlis, D. "Computing protected circumscription" *Journal of Logic Programming*, 4, 1985, pp. 235-249.

(3)   Perlis, D. and Minker, J. "Completeness results for circumscription" *Artificial Intelligence*, 28, 1986, pp. 29-42.

(4)   Drapkin, J. and Perlis, D. "Step-logics: an alternative approach to limited reasoning", *Proceedings of the 7th European Conference on Artificial Intelligence*, Brighton, England, July 1986.

(5)   Minker, J., Perlis, D., and Subramanian, K. "A parallel self-modifying default reasoning system", *Proceedings of the AAAI-86*, Philadelphia, Pennsylvania, August 1986.

(6)   Perlis, D. "Self-Reference, knowledge, belief, and modality", *Proceedings of the AAAI-86*, Philadelphia, Pennsylvania, August 1986.

(7)   Drapkin, J. and Perlis, D. "A preliminary excursion into step logics" *Proceedings of the International Symposium on Methodologies for Intelligent Systems* Knoxville, Tennessee, October 1986.

(8)   Perlis, D. "On the consistency of commonsense reasoning" *Computational Intelligence*, 2, 1986, pp. 180-190.

(9)   Perlis, D. "Commonsense set theory" *Proceedings of Workshop on Meta-Architectures and Reflection* Sardinia, October 1986.

(10) Perlis, D. "Circumscribing with sets" *Artificial Intelligence*, 31, 1987, pp. 201-211.

(11) Perlis, D. "How can a program mean?" *Proc. International Joint Conference on Artificial Intelligence*, Milan, Italy, 1987.

(12) Miller, M. and Perlis, D. "Proving facts about 'I' ". *Proc. International Joint Conference on Artificial Intelligence*, Milan, Italy, 1987.

(13) Gal, A. and Minker, J. "Greater cooperation between database and user: integrity constraints provide an answer" *Proceedings, International Workshop on Natural Language and Logic Programming*, Aug. 17-19, 1987, Vancouver, Canada

(14) Miller, M. and Perlis, D. "Proving self-utterances," *Journal of Automated Reasoning*, pp. 329-338.

(15) Perlis, D. "Circumscription as introspection." *Proceedings, Second International Symposium on Methodologies for Intelligent Systems*, Oct. 14-18, 1987, Charlotte, NC. Z. Ras and M. Zemankova (Eds.), North-Holland.

(16) Perlis, D. "Languages with self-reference II: Knowledge, belief, and modality." *Artificial Intelligence*, 34, 1988, pp. 179-212.

(17) Minker, J. "Perspectives in deductive databases." Invited PODS (Principles of Database Systems conference) lecture, published in *Journal of Logic Programming*, 5, 1988, pp. 33-60.

(18) Lobo, J. and Minker, J. "A meta-program to semantically optimize queries in deductive databases." *Proceedings, Second International Conference on Expert Database Systems*, April 1988, pp. 387-420.

Other references

[A] McCarthy, J. [1980] Circumscription--a form of non-monotonic reasoning. *Artificial Intelligence*, 13 (1,2), pp. 27-39.

[B] McCarthy, J. [1986] Applications of circumscription to formalizing common sense knowledge. *Artificial Intelligence*, 28, pp. 89-116.

[C] Reiter, R. [1982] Circumscription implies predicate completion (sometimes). Proc. AAAI-82, pp. 418-420.

[D] Lifschitz, V. [1985] Computing circumscription. In Proceedings of the Ninth International Joint Conference on Artificial Intelligence, pp. 121-127.

[E] Clark, K. [1978] Negation as failure. In *Logic and Databases*, Gallaire, H. and Minker, J. (eds.) Plenum Press, New York, pp. 293-322.

[F] Etherington, D. W., Mercer, R. E., and Reiter, R. [1985] On the adequacy of predicate circumscription for closed-world reasoning. *Computational Intelligence*, 1, pp. 11-15.

[G] Kueker, D. [1984] Another failure of completeness for circumscription. Week on Logic and Artificial Intelligence, Univ. of Maryland, Oct. 22-26.

[H] Haas, A. [1985] Possible events, actual events, and robots. *Computational Intelligence*, pp.59-70.

[J] Allen, J. [1984] Towards a general theory of action and time. *Artificial Intelligence*, 23, pp.123-154.

[K] McDermott, D. [1982] A temporal logic for reasoning about processes and plans. *Cognitive Science*, 6, pp.101-155.

[L] Sterling, L. and Shapiro, E. [1986] *The Art of Logic Programming*. MIT Press, Cambridge.

[M] Bowen, K. and Kowalski, R. [1982] Amalgamating language and meta-language. In *Logic Programming*, Clark, K. and Tarnlund, S. A. (eds.), Academic Press, London pp. 153-172.

D. LIST OF ALL PARTICIPATING SCIENTIFIC PERSONNEL

Dr. Jack Minker

Dr. Donald Perlis

Dr. Irene Durand (Post Doctoral Visitor)

Ms. Jennifer Elgot-Drapkin

Ms. Teresa Gaasterland

Ms. Annie Gal

Mr. Mark Giuliano

Ms. Diana Gordon (M.S. 1986)

Mr. Hachidai Ito (M.S. 1986)

Dr. Madhur Kohli (Ph.D. 1987)

Ms. Anne Litcher (M.S. 1986)

Mr. Jorge Lobo

Mr. Thomas Melton (M.S. 1988)

Mr. Michael Miller

Dr. Deepak Sherlekar (Ph.D. 1987)

Mr. Krishnan Subramanian (M.S. 1986)